

US007191358B2

(12) **United States Patent**  
**Fujibayashi**

(10) **Patent No.:** **US 7,191,358 B2**  
(45) **Date of Patent:** **Mar. 13, 2007**

(54) **METHOD AND APPARATUS FOR SEAMLESS MANAGEMENT FOR DISASTER RECOVERY**

(75) Inventor: **Akira Fujibayashi**, Sunnyvale, CA (US)

(73) Assignee: **Hitachi, Ltd.**, Tokyo (JP)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

5,943,688 A	8/1999	Fisher et al.
5,999,712 A	12/1999	Moiin et al.
6,035,306 A	3/2000	Lowenthal et al.
6,038,677 A	3/2000	Lawlor et al.
6,061,807 A	5/2000	Albert
6,105,118 A	8/2000	Maddalozzo et al.
6,173,420 B1	1/2001	Sunkara et al.
6,195,732 B1	2/2001	Adams et al.
6,216,211 B1	4/2001	McBrearty et al.
6,269,431 B1	7/2001	Dunham

(Continued)

(21) Appl. No.: **11/228,859**

(22) Filed: **Sep. 16, 2005**

(65) **Prior Publication Data**

US 2006/0041777 A1 Feb. 23, 2006

**Related U.S. Application Data**

(63) Continuation of application No. 10/387,188, filed on Mar. 11, 2003, now abandoned.

(51) **Int. Cl.**  
**G06F 11/00** (2006.01)

(52) **U.S. Cl.** ..... **714/6; 714/5; 707/202**

(58) **Field of Classification Search** ..... None  
See application file for complete search history.

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

4,686,623 A	8/1987	Wallace
5,276,867 A	1/1994	Kenley et al.
5,459,857 A	10/1995	Ludlam et al.
5,511,177 A	4/1996	Kagimasa et al.
5,544,347 A	8/1996	Yanai et al.
5,548,724 A	8/1996	Akizawa et al.
5,737,745 A	4/1998	Matsumoto et al.
5,870,537 A	2/1999	Kern et al.
5,893,919 A	4/1999	Sarkozy et al.
5,933,653 A	8/1999	Ofek

**OTHER PUBLICATIONS**

The Authoritative Dictionary of IEEE Standards Terms, 2000, IEEE Press, Seventh Edition, p. 941.\*

(Continued)

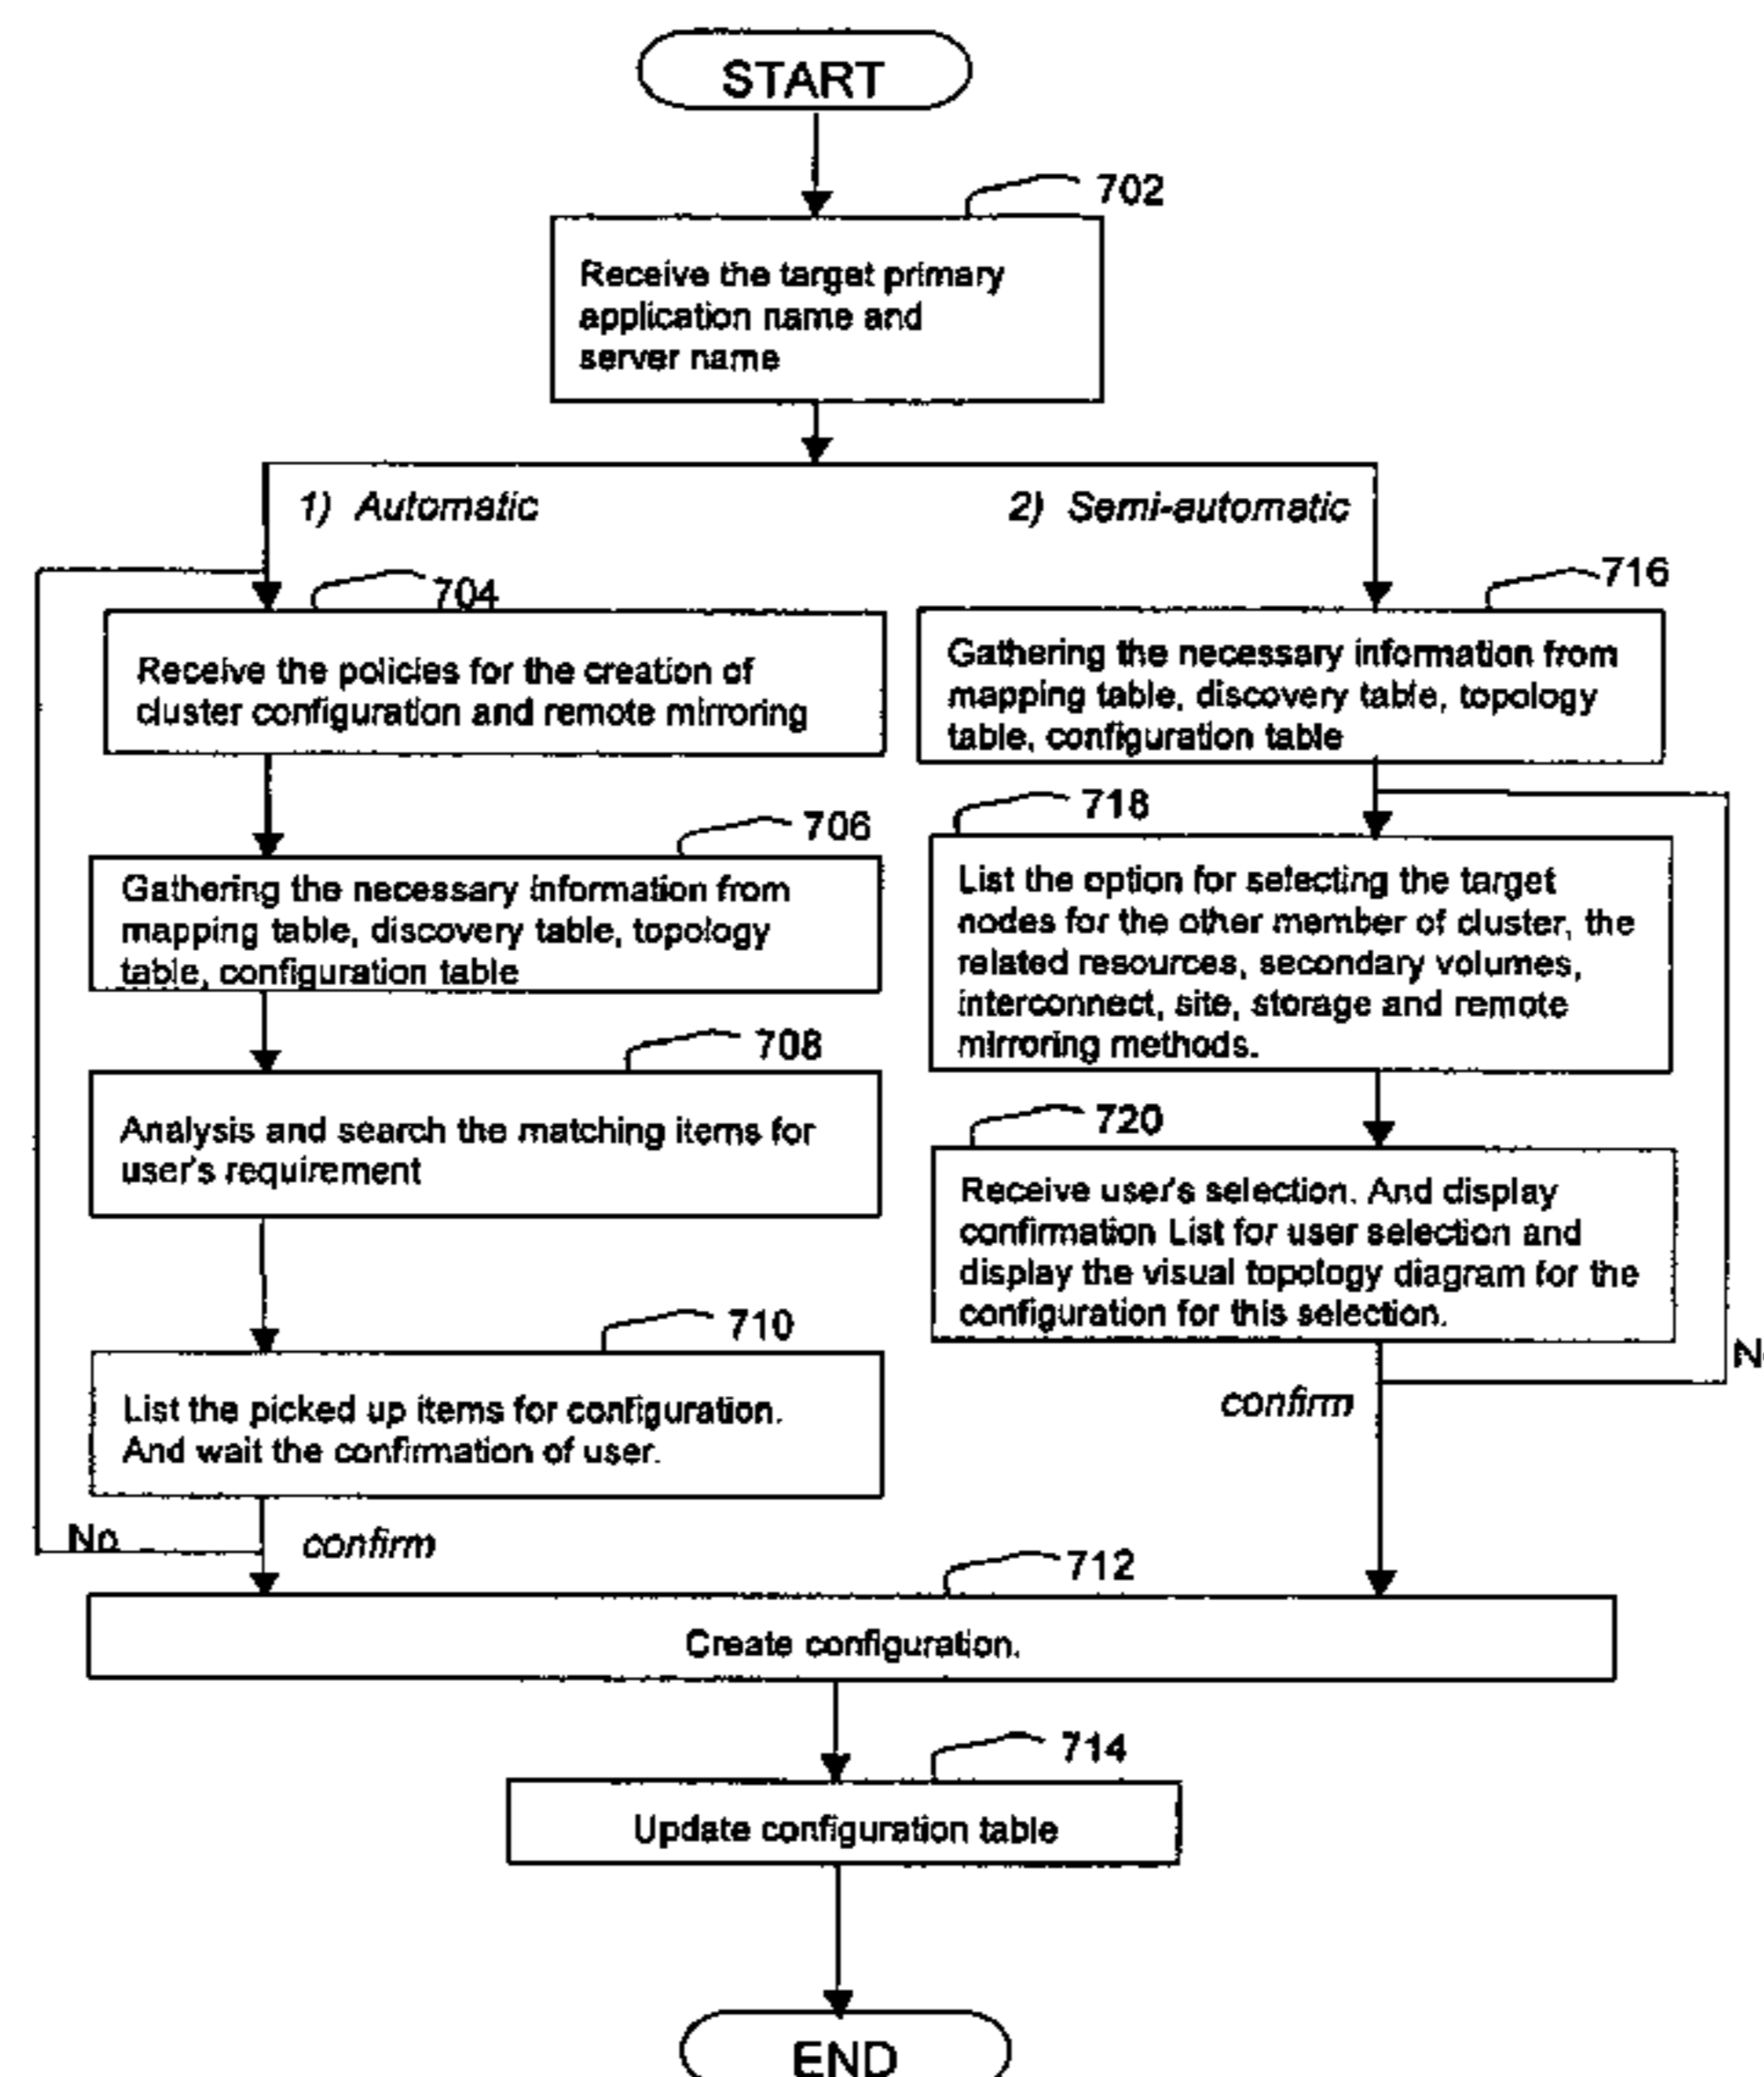
*Primary Examiner*—Marc Duncan

(74) *Attorney, Agent, or Firm*—Townsend and Townsend and Crew LLP

(57) **ABSTRACT**

A method, apparatus, article of manufacture, and system are presented for establishing redundant computer resources. According to one embodiment, in a system including a plurality of processor devices and a plurality of storage devices, the processor devices, the storage devices and the management server being connected via a network, the method comprises storing device information relating to the processor devices and the storage devices and topology information relating to topology of the network, identifying at least one primary computer resource, selecting at least one secondary computer resource suitable to serve as a redundant resource corresponding to the at least one primary computer resource based on the device information and the topology information, and assigning the at least one secondary computer resource as a redundant resource corresponding to the at least one primary computer resource.

**5 Claims, 12 Drawing Sheets**



U.S. PATENT DOCUMENTS

6,324,654 B1 11/2001 Wahl et al.  
6,393,485 B1 5/2002 Chao et al.  
6,425,049 B1 7/2002 Yamamoto et al.  
6,438,705 B1 8/2002 Chao et al.  
6,597,862 B1 7/2003 Saeki et al.  
6,606,643 B1 8/2003 Emens et al.  
6,618,819 B1\* 9/2003 Adamovits et al. .... 714/13  
6,633,955 B1 10/2003 Yin et al.  
6,665,780 B1 12/2003 Bradley  
6,691,245 B1 2/2004 DeKoning et al.  
6,745,341 B1 6/2004 Onitsuka et al.  
6,785,678 B2 8/2004 Price  
2001/0044807 A1 11/2001 Kleiman et al.  
2002/0004857 A1 1/2002 Arakawa et al.  
2002/0007445 A1 1/2002 Blumenau et al.  
2003/0046602 A1 3/2003 Hino et al.

2003/0233518 A1 12/2003 Yamagami et al.  
2004/0019822 A1 1/2004 Knapp  
2004/0030851 A1 2/2004 Ohran et al.  
2004/0078397 A1 4/2004 Mehta et al.  
2004/0098637 A1 5/2004 Duncan et al.  
2004/0230863 A1 11/2004 Buchhorn

OTHER PUBLICATIONS

Clarke "Technical Introduction to Oracle Fail Safe, A Guide to Concepts and Terminology," An Oracle Technical White Paper (2001).  
"Managing Resource Definitions in Distributed Systems," IBM Technical Disclosure Bulletin pp. 81-84 (Dec. 1995).  
"Methodology for Profile Replication," IBM Technical Disclosure Bulletin (Nov. 1992).

\* cited by examiner

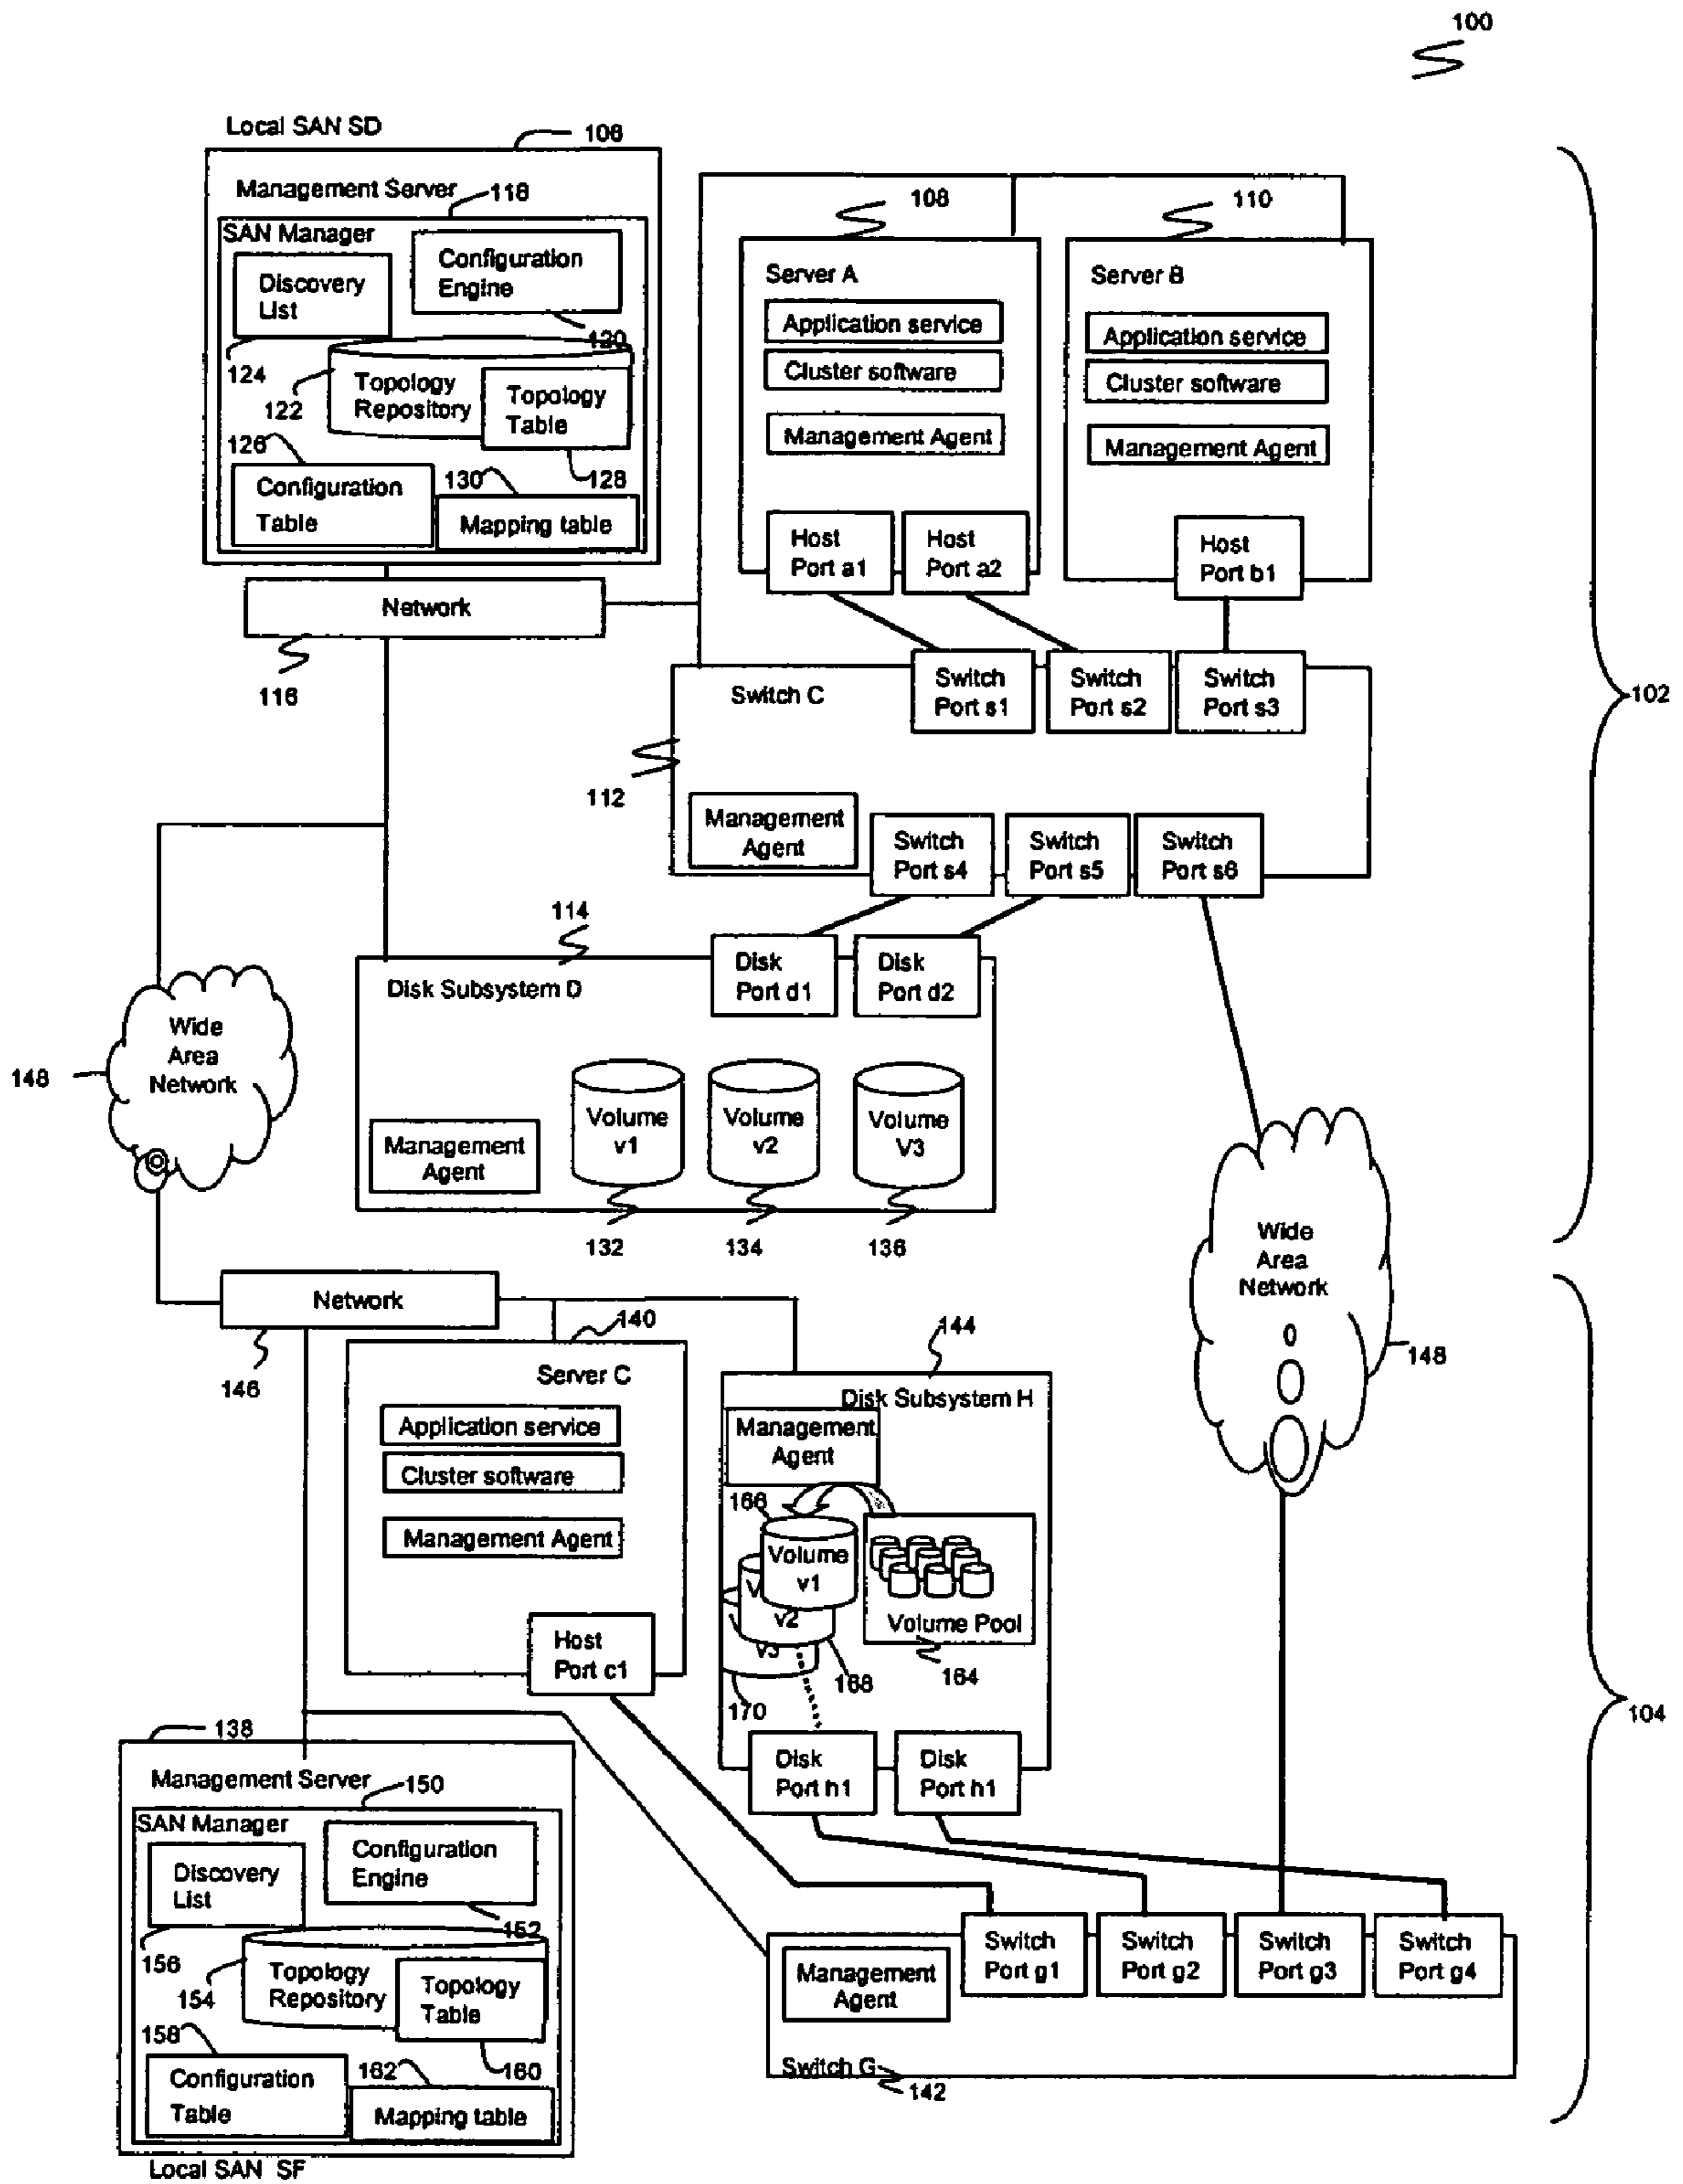


Fig. 1

Σ<sup>130</sup>

ID	Server	Application	Related mount point	Related Volume ID	DG ID	Block device	LU Binding ID	SCSI ID	SCSI LUN
1	Server A	DB	/u01	Vol1	0	c2t2d1	1	2	1
			/u02	Vol2	1	c2t3d1	2	3	1
			/u03	Vol3	2	c3t4d2	3	4	2
2	Server A	File service	/u05	-	-	c3t4d3	4	4	3
			/u06	-	-	c3t4d4	5	4	4
...	...	...	...	...	...	...	...	...	...

Σ 202   Σ 204   Σ 206   Σ 208   Σ 210   Σ 212   Σ 214   Σ 216   Σ 218   Σ 220

Fig. 2

300

Binding ID	SCSI ID	LUN	Inquiry Information
1	2	1	Vendor D, Storage D, Volume ID 1
2	3	1	Vendor D, Storage D, Volume ID 2
3	4	2	Vendor D, Storage D, Volume ID 2

302      304      306      308

Fig. 3

§ 124

Local SAN ID	Discovery ID	Device Type	Device Information	IP Address	Area/Global Position
1	1	Server	Vendor A, Server A	100.100.100.100	SD Latitude xxxxN Longitude xxxxW
1	2	Server	Vendor B, Server B	100.100.100.101	SD Latitude xxxxN Longitude xxxxW
1	3	FC Switch	Vendor C, Switch C	100.100.100.102	SD Latitude xxxxN Longitude xxxxW
1	4	Disk Subsystem	Vendor D, Storage D	100.100.100.103	SD Latitude xxxxN Longitude xxxxW
2	1	Server	Vendor A, Server C	120.100.100.105	SF Latitude xxxxN Longitude xxxxW
2	2	Disk Subsystem	Vendor D, Storage H	120.100.100.106	SF Latitude xxxxN Longitude xxxxW
2	3	FC Switch	Vendor C, Switch G	120.100.100.107	SF Latitude xxxxN Longitude xxxxW

402      404      406      408      410      412

Fig. 4A

440

Local SAN ID	Discovery ID	Function Type	Device Information
1	1	Host-based Remote mirroring	Vendor A, Server A, enterprise-class
1	2	Storage-based Remote Mirroring	Vendor B, Storage B, mid-range
1	3	Storage-based Remote Mirroring	Vendor C, Storage C, mid-range
1	4	Storage-base Remote Mirroring	Vendor D, Storage D, enterprise-class
2	1	Storage-base Remote Mirroring	Vendor D, Storage H, enterprise-class
2	2	Host-based Remote Mirroring	Vendor K, server C, enterprise-class

442

444

446

448

Fig. 4B

500  
S

502 Server			504 Local network				506 Interconnect	508 Local network			510 Storage			
Server Name	Binding ID	Host Port ID	Switch Name	Switch Port ID	Switch Name	Switch Port ID		Switch Name	Switch Port ID	Switch Name	Switch Port ID	Storage Name	Disk Port ID	Volume ID
Server A	LU1	a1	SW C	s1	SW C	s4						Storage D	d1	V1
Server A	LU2	a2	SW C	s2	SW C	s4						Storage D	d1	V2
Server A	LU3	a2	SW C	s2	SW C	s4						Storage D	d1	V3
Server B	LU1	b1	SW C	s3	SW C	s4						Storage D	d1	V1
Server C	LU1	c1	SW G	g1	SW G	g2						Storage H	h1	V1
					SW C	s6	SF-SD	SW G	g3					
								SW C	s6	SW C	s5	Storage D	d2	
								SW G	g3	SW G	g4	Storage D	d1	

Fig. 5



Zone ID	Switch Port ID List
1	(s1,s2,s3,s4)
2	(s5,s6)

602

604

600

Fig. 6A

Zone ID	Switch Port ID List
1	(g1, g2)
2	(g3, g4)

602

604

606

Fig. 6B

608

610	612	614	616			618			620
ID	Group ID	Group name	Primary host list	Primary volume ID	Primary Volume Subsystem ID	Secondary host list	Secondary Volume#1 Volume ID	Secondary Volume#1 Subsystem ID	Cluster Config ID
1	1	OraDB	Host A / xxx.xxx.xxx.xxx	1	2022345	Host C / xxx.xxx.xxx.xxx	1	3004567	1
				2	2022345		2	3004567	1
				3	2022345		3	3004567	1
2	...	...	...	...	...	...	...	...	...

Fig. 6C

622

624	626	628	630	632					634			636				
ID	valid	Grp ID	Grp Name	Primary host list	Primary Logical Volume ID	Primary DG ID	Primary Block Device	Primary Vol ID	Primary Subsystem ID	Secondary host list	Secondary Logical Volume ID# / name	Secondary DG ID# / name	Secondary Block Device	Secondary Volume#1 Volume ID	Secondary Volume#1 Subsystem ID	Cluster Config ID
1	1	1	FSGF	Host B / xxx.xx x.xxx xxx	1 / vol1	1 / DG1	C210d0	1	2022345	Host C / xxx.xxx.xx x.xxx	2 / vol2	1 / DG1	C210d3	3	3004567	1
2	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...

Fig. 6D

640 ID	642 valid	644 Cluster ID / name	646 Cluster type / vendor	648 Member node list	650 Heartbeat list	652 Heartbeat configuration ID list	654 Replication type list	656 Replication configuration ID list
1	1	1 / Clus1	XX Cluster server / Vendor A	Host A / xxx.xxx.xxx.xxx Host C / xxx.xxx.xxx.xxx	1 / Private-LAN 2 / Storage-based HB	1 2	Host-based Storage-based	1
2	0	2 / Clus2	YY Cluster server / Vendor C	Host B / xxx.xxx.xxx.xxx Host D / xxx.xxx.xxx.xxx	1 / Private-LAN	Host-based	Host-based	2
3	1	...	...	...	...	...	...	...

Fig. 6E

660 ID	662 valid	664 Cluster type ID	666 Resource Grp ID	668 Resource Grp Name	670 Member node list	672 Resource List	674 Replication Type	676 Replication Config ID
1	1	1 / xxx cluster server / vendor A	1	DB_Grp	Host A / xxx.xxx.xxx.xxx Host C / xxx.xxx.xxx.xxx	DB NIC IP FS DG Replica	Storage-based	1
2	0	2 / yyy cluster server / vendor C	2	FS_Grp	Host B / xxx.xxx.xxx.xxx Host C / xxx.xxx.xxx.xxx	NFS NIC IP FS DG Replica	Host-based	2
3	1	...	...	...	...	...	...	...

Fig. 6F

678

680	682	684	686	688	690	692	694
ID	valid	Cluster type ID	Heartbeat type ID	Heartbeat Name	Member node list	NIC List	Storage list
1	1	1 / xxx cluster server / vendor A	1	HB-storage	Host A / xxx.xxx.xxx.xxx Host C / xxx.xxx.xxx.xxx	-	3002345 2000456
2	0	2 / yyy cluster server / vendor C	2	HB-Host	Host B / xxx.xxx.xxx.xxx Host D / xxx.xxx.xxx.xxx	LAN-connection #1	-
3	1	1 / xxx cluster server / vendor A	2	HB-Host	Host A / xxx.xxx.xxx.xxx Host C / xxx.xxx.xxx.xxx	LAN-connection #1 LAN-connection #1	-
4	1	...	...	...	...	...	...

Fig. 6G




Fig. 7




Fig. 8

## METHOD AND APPARATUS FOR SEAMLESS MANAGEMENT FOR DISASTER RECOVERY

### CROSS-REFERENCES TO RELATED APPLICATIONS

This application is a continuation of U.S. patent application Ser. No. 10/387,188, filed Mar. 11, 2003, now abandoned, the entire disclosure of which is incorporated herein by reference.

### BACKGROUND OF THE INVENTION

This invention is generally related to the field of clustering systems and remote mirroring technology.

The use of clustering systems to accomplish fault-tolerance and/or load-balancing is becoming increasingly popular. Generally speaking, a clustering system may provide redundant resources so that if one portion of the system experiences failure, another portion can take over affected tasks or otherwise provide recovery from the failure. Also, a clustering system may use its redundant resources to process tasks in a more distributed manner, allowing different portions of the system to work in parallel in accomplishing tasks.

A typical clustering system may be made up of two or more nodes, each having its own processing and storage capabilities. In one particular use of a clustering system, a primary node may comprise of a server and associated storage devices, while a secondary node may also comprise of another server and associated storage devices. The secondary node may be created to be similar to the primary node, in terms of processing, storage, and other capabilities. Here, the clustering system may maintain exact correspondence between the data storage of the primary node and the data storage of the secondary node, such that any write or read to data storage at the primary node is replicated at the secondary node. If the primary node fails as it performs its various tasks, the secondary node may take over the tasks performed by the primary node. For example, if a web server that is configured as a primary node in a clustering system fails for some reason, a secondary node may take over and serve web server functions in place of the failed primary node. A web site supported by such a system thus continues to operate with little or no down time. Web site visitors may continue to visit the associated web site as if no failure had occurred. In this example, providing a primary and a secondary node of similar capabilities allows the secondary node to be capable of taking over the tasks previously performed by the primary node.

In other situations, the secondary node may have lesser capabilities than the primary node. For example, if the secondary node is only designed to temporarily take over the tasks of the primary node, or if the secondary node is only designed to record periodic snapshots of the data storage of the primary node, it may be sufficient to create the secondary node with lesser capabilities. This may be especially true if the cost associated with creating a similarly capable secondary node is to be avoided, or if failure of the primary node is not expected to extend beyond a certain amount of time. Thus, depending on the situation, the required capabilities of the secondary node may vary.

The correspondence between the data storage of a primary node and the data storage of a secondary node storage may also be referred to as remote mirroring. This is especially the case if the data storage of the primary node is at a geographically distant location from the data storage of the

secondary node. Remote mirroring may be carried out by different portions of a system. For example, in host-based remote mirroring, a host, such as a server, may be principally responsible for maintaining the correspondence between the data storage of the primary node and the data storage of the secondary node. In storage-based remote mirroring, a storage system, such as a storage area network (SAN), may be principally responsible for maintaining such correspondence. Depending on the implementation, remote mirroring may require separate software and equipment installation and/or configuration, in addition to that required by other parts of the clustering system.

Currently, in order to realize the many advantages of a clustering system, the multiple nodes of a clustering system must be established by a system administrator. For example, in a clustering system having a primary and a secondary node, the system administrator must decide exactly what should be the processing, storage, and other capabilities of the secondary node, install or identify available resources meeting those capabilities, install required software, and perform necessary configurations to set up the clustering system. These steps involve factors that can be overwhelmingly complex and difficult to analyze for the system administrator, even if that person is an expert. Thus, the administrator may only be able to make a rough guess, in an ad hoc manner, as to what storage capability is needed for the secondary node. As discussed above, the required storage capability of the secondary node may vary from situation to situation, and it may not always be ideal to simply mimic the storage capability of the primary node.

Furthermore, after the desired processing, storage, and other capabilities of the secondary node is decided, the administrator must go about looking for existing equipment in the system that fit the description, or install such equipment. In a large system having many different components, it may be extremely difficult and time-consuming for an administrator to search through all available resources in order to find the appropriate equipment. Finally, after the appropriate resources are decided and located, software installation and configuration may take additional time and effort. Thus, while clustering systems provide import fault-tolerance and/or load-balancing capabilities, the deployment of clustering systems remains largely a difficult and imprecise undertaking.

### BRIEF SUMMARY OF THE INVENTION

The present invention provides a method, apparatus, article of manufacture, and system for establishing redundant computer resources. According to one embodiment, in a system including a plurality of processor, a plurality of storage devices, and a management server connected via a network, the method comprises storing device information relating to the processor devices and the storage devices and topology information relating to topology of the network, identifying at least one primary computer resource, the at least one primary computer resource including at least one primary processor device and at least one portion of storage implemented in at least one primary storage device, selecting at least one secondary computer resource suitable to serve as a redundant resource corresponding to the at least one primary computer resource based on the device information and the topology information, the at least one secondary computer resource including at least one secondary processor device and at least one portion of storage implemented in at least one secondary storage device, and assign-

ing the at least one secondary computer resource as a redundant resource corresponding to the at least one primary computer resource.

If the at least one primary storage device has storage-based remote mirroring function, the at least one secondary computer resource may be selected such that the at least one secondary storage device also has storage-based remote mirroring function and is accessible from the at least one primary storage device.

In one embodiment, the at least one secondary computer resource is selected based on at least one user-specified policy, which may include performance of the at least one secondary computer resource, reliability of the at least one secondary computer resource, and/or cost of the at least one secondary computer resource.

In another embodiment, the step for selecting the at least one secondary computer resource comprises the steps of selecting at least one candidate suitable to serve as a redundant resource corresponding to the at least one primary computer resource, presenting the at least one candidate to a user, and receiving input from the user indicating selection, from the at least one candidate, of the at least one secondary computer resource.

#### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a clustering system in accordance with at least one embodiment of the present invention.

FIG. 2 is an illustration of a mapping table.

FIG. 3 is an illustration of a logical unit number (LUN) binding table.

FIG. 4A is an illustration of a discovery list.

FIG. 4B is an illustration of a functional discovery list that may be maintained in addition to or in place of the discovery list shown in FIG. 4A.

FIG. 5 is an illustration of a topology table.

FIG. 6A illustrates a fibre channel switch (FC-SW) zoning configuration table.

FIG. 6B illustrates a different FC-SW zoning configuration table.

FIG. 6C illustrates a storage-based replication configuration table.

FIG. 6D illustrates a host-based replication configuration table.

FIG. 6E illustrates a cluster configuration table.

FIG. 6F illustrates a cluster resource group configuration table.

FIG. 6G illustrates a heartbeat configuration table.

FIG. 7 is a flow chart summarizing the general steps involved in automatic configuration and semi-automatic configuration of a clustering system in accordance with at least one embodiment of the present invention.

FIG. 8 depicts a visual configuration diagram that may be presented to the user.

#### DETAILED DESCRIPTION OF THE INVENTION

##### Clustering System

FIG. 1 is a block diagram of a clustering system 100 in accordance with at least one embodiment of the present invention. Here, clustering system 100 is comprised of equipment found in at least two geographically distinct locations 102 and 104. For example, location 102 may be a metropolitan area such as San Diego, Calif., and location

104 may be a different metropolitan area such as San Francisco, Calif. At location 102, a management server 106 is responsible for monitoring, configuring, and otherwise managing servers 108 and 110, network equipment 112, and storage equipment 113, 114, and 115. Management server 106, servers 108 and 110, network equipment 112, and storage equipment 113, 114, and 114 communicate through a local network 116, forming a local SAN.

As shown, management server 106 includes a SAN manager 118 that includes a configuration engine 120 and a topology repository 122. SAN manager 118 also maintains a discovery list 124, a configuration table 126, a topology table 128, and a mapping table 130, which are discussed in further detail below. SAN manager 118 maintains this information by communicating with various management agents located in servers 108 and 110, network equipment 112, and storage equipment 113, 114, and 115. SAN manager 118 and the various management agents may be implemented in software.

Server 108 may include one or more application programs. These application programs may be server level applications such as Web server applications, network file sharing applications, and others. As FIG. 1 illustrates, server 108 may also include clustering software for maintaining a clustering system, a management agent, and a number of host ports. Server 110 is similarly arranged and may also include one or more application programs, clustering software, a management agent, and a number of host ports.

Network equipment 112 is illustrated in FIG. 1 as a switch having a number of switch ports. Network equipment 112 also includes a management agent. Network equipment 112 facilitates communication through local network 116. As shown, network equipment 112 provides communication between servers 108 and 110 and storage equipment 115.

Storage equipment 115 may include a number of disk ports, a number of logical volumes 132, 134, and 136, and a management agent. Here, each of the logical volumes 132, 134, and 136 may be implemented in different ways, such as by use of various types of redundant array of independent disks (RAID). Each of logical volumes 132, 134, 136 may be implemented on a single physical disk (not shown), across multiple physical disks (not shown) within a disk group (not shown), across disks in multiple disk groups, or in some other arrangement.

Here, server 108, network equipment 112, and storage equipment 115 may represent a primary node in a clustering system. For example, server 108 may be executing a database application, using storage equipment 115 to store the associated databases and communicating data to and from storage equipment 115 through network equipment 112. Fault-tolerance for this database service may be realized by creating a secondary node corresponding to the primary node. Use of equipment located at a geographically distinct location, such as location 104, would provide effective fault-tolerance because if a catastrophic local event damages equipment at location 102, redundant equipment at location 104 would be able to provide effective recovery.

At location 104, a management server 138 is responsible for monitoring, configuring, and otherwise managing a server 140, network equipment 142, and storage equipment 144. Management server 138, server 140, network equipment 142, and storage equipment 144 communicated through a local network 146, forming a local SAN. Local SANs at locations 102 and 104, and perhaps other local SANs, may together form a wide area SAN by communicating over one or more wide area networks 148.



As shown, management server **138** includes a SAN manager **150** that includes a configuration engine **152** and a topology repository **154**. SAN manager **150** also maintains a discovery list **156**, a configuration table **158**, a topology table **160**, and a mapping table **162**, which are discussed in further detail below. SAN manager **150** maintains this information by communicating with various management agents located in server **140**, network equipment **142**, and storage equipment **144**. SAN manager **150** and the various management agents may be implemented in software.

Server **140** may include one or more application programs, clustering software for maintaining a clustering system, a management agent, and a number of host ports. Network equipment **142** is illustrated in FIG. 1 as a switch having a number of switch ports. Network equipment **142** also includes a management agent. Network equipment **142** facilitates communication through local network **146**. As shown, network equipment **142** provides communication between server **140** and storage equipment **144**.

Storage equipment **144** may include a number of disk ports, a pool **164** of logical volumes, from which logical volumes **166**, **168**, and **170** may be selected, and a management agent. Here, each of the logical volumes in logical volume pool **164**, including logical volumes **166**, **168**, and **170**, may be implemented in different ways, such as by use of various types of redundant array of independent disks (RAID). Thus, each of the logical volumes may be implemented on a single physical disk (not shown), across multiple physical disks (not shown) within a disk group (not shown), across disks in multiple disk groups, or in some other arrangement.

Here, server **140**, network equipment **142**, and storage equipment **144** may be used to form a secondary node associated with the previously discussed primary node in the clustering system. For example, if the clustering system is designed to provide a secondary node having similar processing, storage, and other capabilities as those of the primary node, it would be desirable to identify a secondary node having similar equipment as the primary node. Server **140**, network equipment **142**, and storage equipment **144** may fit such requirements. The present invention allows equipment such as server **140**, network equipment **142**, and storage equipment **144** to be identified as resources that may be used to form the secondary node.

Servers **108**, **110**, and **140** are examples of processor devices, network equipment **115** and **144** are examples of storage devices, and network equipment **112** and **142** are examples of network interface devices.

Information Maintained at Management Server and Elsewhere

FIG. 2 is an illustration of mapping table **130** maintained in management server **106** of FIG. 1. Mapping table **130** is illustrated here as an example. Other mapping tables, such as mapping table **162** maintained in management server **138**, may have similar formats. As shown in FIG. 2, mapping table **130** provides a mapping between application programs being executed and the location(s) of data storage being utilized by such application programs. For instance, an application program executing in server **108** may utilize logical volumes **132**, **134**, and **136** in storage equipment **115**, and mapping table **130** would register such utilization in detail. Different methods may be used to identify the various application programs executing in a particular server. One such method involves using the Common Information Model (CIM) standard, which allows application programs executing in a server may communicate with one another.

For example, the management agent in server **108** may use the CIM standard to communicate with, and thereby identify, the various application programs executing in server **108**. Another method involves using repository information maintained by the operating system of the server. For example, the management agent in server **108** may retrieve data from the repository information of the operating system of server **108** to identify various application program executing in server **108**.

Mapping table **130** is shown to include the following categories of information: ID **202**, Server **204**, Application **206**, Related Mount Point **208**, Related Volume ID **210**, Disk Group (DG) ID **212**, Block Device **214**, Logical Unit (LU) Binding ID **216**, Small Computer System Interface (SCSI) ID **218**, and SCSI Logical Unit Number (LUN) **220**. Here, table **130** indicates that a database (DB) application is executing in Server A (server **108**). Table **130** further indicates that this DB application is utilizing logical volumes Vol1, Vol2, and Vol3 (logical volumes **132**, **134**, and **136**). For each of these three logical volumes, table **130** provides additional information. Taking Vol1 just as an example, table **130** indicates the mount point (**/u01**) at which Vol1 is associated with, or “mounted” to, the system executing the DB application. Table **130** also indicates the physical disk group (**0**) and block device (**c2t2d1**) in which Vol1 is implemented. In this example, logical volumes are also associated with SCSI IDs, as well as LUNs within particular SCSI IDs. Here, Vol1 is shown to be associate with a particular SCSI ID (**2**) and a particular SCSI LUN (**1**).

FIG. 3 is an illustration of a LUN binding table **300** maintained in server **108** of FIG. 1. LUN binding table **300** is illustrated here as an example. Other LUN binding tables maintained in other servers, such as servers **110** and **140**, may have similar formats. LUN binding table **300** indicates the SCSI ID assignment and LUN assignment associated with location(s) of data storage being utilized by application programs executing in server **108**. LUN binding table **300** is shown to include the following categories of information: Binding ID **302**, SCSI ID **304**, LUN **306**, and Inquiry Information **308**. Each Binding ID **302** indicates a particular location of storage and is associated with a particular SCSI ID **304** and a particular LUN **306**. Also, each Binding ID **302** further indicates Inquiry Information **308**, which can provide additional data such as vendor, storage type, and logical volume information. Binding table **300** may be maintained as a part of the operation of the management agent in server **108**. Thus, individual binding tables maintained at various servers, such as servers **108** and **110**, may be used to form the mapping table **130** shown in FIG. 2.

FIG. 4A is an illustration of discovery list **124** maintained in management server **106** of FIG. 1. Discovery list **124** is illustrated here as an example. Other discovery lists, such as discovery list **156** maintained in management server **138**, may have similar formats. As shown in FIG. 4, discovery list **124** provides a listing of devices available at various locations, such as locations **102** and **104**. Discovery list **124** shows the following categories of information for each device: Local SAN ID **402**, Discovery ID **404**, Device Type **406**, Device Information **408**, IP address **410**, and Area/Global Position **412**. Local SAN ID **402** identifies the local SAN to which the device belongs. Discovery ID **404** identifies a numerical order for the device within its local SAN. Device Information **408** may indicate various information relating to the device, such as vendor and device type. IP address **410** indicates the IP address assigned to the device. Area/Global Position **412** provides information relating to the location of the device, such as name of metropolitan

area, longitude, and latitude. Thus, discovery list 124 allows management server 106 to identify available devices at various locations, including distant locations, that may be potential resources suitable to serve as part of a secondary node corresponding a primary node in a clustering system.

FIG. 4B is an illustration of a functional discovery list 440 that may be maintained in management server 106 of FIG. 1, in addition to or in place of discovery list 124. Functional discovery list 440 is illustrated here as an example. Other discovery lists maintained in other management servers may have similar formats. As shown in FIG. 4B, functional discovery list 440 provides a listing of devices available at various locations, such as locations 102 and 104. Functional discovery list 440 shows the following categories of information for each device: Local SAN ID 442, Discovery ID 444, Function Type 446, and Device Information 448. Local SAN ID 442 identifies the local SAN to which the device belongs. Discovery ID 444 identifies a numerical order for the device within its local SAN. Function Type 446 provides information on the possible function of the device, such as use in host-based remote mirroring or storage-based remote mirroring. Device Information 448 may indicate various information relating to the device, such as vendor, device type, and device class. Functional discovery list 440 allows management server 106 to identify available devices at various locations, including distant locations, that may be potential resources suitable to serve as part of a secondary node corresponding a primary node in a clustering system.

FIG. 5 is an illustration of topology table 128 maintained in management server 106 of FIG. 1. Topology table 128 is illustrated here as an example. Other topology tables, such as topology table 160 maintained in management server 138, may have similar formats. As shown in FIG. 5, topology table 128 provides a summary of interconnections over which data may be sent in system 100. Topology table 128 shows the following categories of information: server information 502, first local network information 504, interconnect information 506, second local network information 508, and storage information 510. Topology table 128 depicts the manner by which various networking and storage equipment are linked, including local and wide area network connections. Here, topology table 128 is shown to be focused on storage network topology for purposes of illustration. Other types of topology information may be included as well.

FIGS. 6A–6G show various configuration tables that may be implemented, individually or in combination, as the contents of configuration table 126 maintained in management server 106 of FIG. 1. Contents of configuration table 126 is illustrated here as examples. Other configuration tables, such as configuration table 158 maintained in management server 138, may have similar formats.

FIG. 6A illustrates a fibre channel switch (FC-SW) zoning configuration table 600. This table contains categories of information including Zone ID 602 and Switch Port ID List 604. Zone ID 602 identifies different zones, or groupings of devices, such that devices within a common zone may readily communicate with one another. Switch Port ID List 604 identifies the different network ports which belong to the identified zone. FIG. 6B illustrates a different FC-SW zoning configuration table 606, similar in structure to table 600. Zoning configuration tables 600 and 606 allow convenient separation of groups of devices. Here, tables 600 and 606 are described as fibre channel switch zoning configuration tables for purposes of illustration, other types of equipment may also be organized in similar zoning tables.

FIG. 6C illustrates a storage-based replication configuration table 608. This table identifies the configuration of

storage-based data replication from a set of primary storage locations to a corresponding set of secondary storage locations. Here, the storage system is responsible of maintaining the proper replication of data. Table 608 shows the following categories of information: ID 610, Group ID 612, Group Name 614, primary storage information 616, secondary storage information 618, and Cluster Config ID 620. ID 610 is an entry identifier. Group ID 612 and Group Name 614 relate to the identification number and name for each group of storage resources, such as a group of volumes, representing a storage location. The primary and secondary storage information 616 and 618 each identifies the host and volume information associated with the relevant storage location. Cluster Config ID 620 identifies a label for the cluster corresponding to the primary and secondary storage locations.

FIG. 6D illustrates a host-based replication configuration table 622. This table identifies the configuration of host-based data replication from a set of primary storage locations to a corresponding set of secondary storage locations. Here, the host system is responsible of maintaining the proper replication of data. Table 622 shows the following categories of information: ID 624, Valid 626, Group ID 628, Group Name 630, primary storage location information 632, secondary storage location information 634, and Cluster Config ID 636. Valid 626 relates to whether the particular replication configuration is available. Also, primary and secondary storage location information 632 and 634 are each shown to also include information for identifying the corresponding disk group and block device. Other information in table 622 is similar to information shown in table 608 of FIG. 6C.

FIG. 6E illustrates a cluster configuration table 638. This table identifies the arrangement of various clusters in the system, which may include the configuration of physical devices being controlled by cluster software. Table 638 shows the following categories of information: ID 640, Valid 642, Cluster ID/Name 644, Cluster Type/Vender 646, Member Node List 648, Heartbeat List 650, Heartbeat Configuration ID List 652, Replication Type List 654, and Replication Configuration ID List 656. ID 640 identifies a numeric label for each entry, Valid 642 relates to whether the particular cluster is available. Cluster ID/Name 644 provides a number identifier and a name identifier for each cluster presented. Cluster Type/Vendor 646 identifies the classification of the cluster and vendor of the associated equipment. Member Node List 648 identifies the nodes that are members of the particular cluster. Heartbeat List 650 and Heartbeat Configuration 652 relate to arrangement of the heartbeat, which provides a signal that may be used to indicate whether a node, or particular resource at a node, is active. Replication Type List 654 and Replication Configuration IID List 656 relate to the type of replication available and the associated configuration label.

FIG. 6F illustrates a cluster resource group configuration table 658. This table identifies the various resources available at different clusters, which may include the configuration of the logical resource group for each node in each cluster. Such resources may be processing, communication, storage, or other types of resources. Table 658 shows the following categories of information: ID 660, Valid 662, Cluster Type ID 664, Resource Group ID 666, Resource Group Name 668, Member Node List 670, Resource List 672, Replication Type 674, and Replication Configuration ID 676. ID 660 provides an numerical label for each entry, Valid 662 relates to whether the particular cluster is available. Cluster Type ID 664 provides an identifier for the

cluster and indicates the type and vendor of equipment associated with the cluster. Resource Group ID **666** and Resource Group Name **668** provide a number identifier and a name identifier for each collection of resources associated with the cluster. Resource List **672** identifies the particular resources available within the identified resource group. Replication Type **674** and Replication Config iD **676** relate to the type of replication available and the associated configuration label.

FIG. **6G** illustrates a heartbeat configuration table **678**. This table identifies provides further detail on the arrangement of the heartbeat for each cluster. Table **678** shows the following categories of information: ID **680**, Valid **682**, Cluster Type ID **684**, Heartbeat Type ID **686**, Heartbeat Name **688**, Member Node List **690**, NIC List **692**, and Storage List **694**. ID **680** provides a numerical label for each entry. Valid **682** relates to whether the cluster is available. Cluster Type ID **684** provides an identifier for the cluster and indicates the type and vendor of equipment associated with the cluster. Heartbeat Type ID **686** and HeartBeat Name **688** identify the classification and name of the heartbeat utilized. For example, the heartbeat may be host-based or storage-based. Member Node List **690** identifies the nodes that are members of the particular cluster. NIC List **692** identifies NICs which correspond the to a particular host-base heartbeat. Storage list identifies storage systems which correspond to a particular storage-based heartbeat.

The information maintained at each management server may be communicated to other management servers. For example, although management servers **106** and **108** are situated at geologically distinct locations **102** and **104**, respectively, they may exchange some or all of the information that is contained in various tables such as those discussed above.

#### Automatic Configuration

FIG. **7** is a flow chart summarizing the general steps involved in automatic configuration and semi-automatic configuration of a clustering system in accordance with at least one embodiment of the present invention. The steps shown may be implemented as an integrated routine that allows the selection of either automatic configuration or semi-automatic configuration. Alternatively, the steps shown may be implemented as two separate routines. That is, a system may employ only automatic configuration, or only semi-automatic configuration. For purposes of illustration, FIG. **7** shows the establishment of a clustering system through the formation of a secondary node corresponding to a primary node. Different steps shown in FIG. **7** may be accomplished with use of a user interface, such as an interactive graphical user interface (GUI). Also, the GUI can be situated at any location, as long as the relevant information can be passed to the system. For example, the information submitted through the GUI by the user may be sent to the management server **106**, or to the management server **138**.

Under automatic configuration, establishment of a clustering system begins with step **702**, in which the primary node of the planned clustering system is identified. This may involve identification, by the user, of the name of one or more target applications and the name of the target server corresponding to the primary node. Alternatively, a more automated process may be employed. For example, the main application executing in a target server may be selected.

Next, in step **704**, policies for creating the clustering system, including remote mirroring features, may be specified. This step may involve specification by the user of

general policies to follow in establishing the clustering system and importance assigned to such policies. For example, the user may be presented with three potential policies: (1) performance, (2) reliability, and (3) cost.

Performance may relate to the effectiveness of the data transfer between the data storage of the primary node and the data storage of the secondary node, which may involve measures of bandwidth, distance, and network usage in a wide area SAN covering metropolitan areas of San Francisco (SF) and San Diego (SD) are provided in the table below:

		Network	
	Type	Total	Usage
SD	Local	2 Gbps	50%
SF-SD	Interconnect	48 Gbps	10%
SF	Local	2 Gbps	8%

Illustrative measures of bandwidth, distance, and network usage in the same wide area SAN, but from the perspective of the San Diego (SD) metropolitan area, are provided in the table below:

		Tested		Network	
	Type	Throughput	Distance	Total	Usage
SF	interconnect	500 Mbps	1000 mile	48 Gbps	10%

Thus, if a user places emphasis on performance, the secondary node may be chosen to have equal performance as the primary node, in terms of processing capability (server type), storage capability (throughput, cache size, RAID level, etc.), and network interface capability (number and performance of host bus adaptors). Also if there are two or more option for interconnects between the primary device and secondary device, the interconnect that has more available throughput capacity may be chosen. For example, assume there are two interconnects: interconnect A, which has 48 Gbps total throughput capacity and 10% average usage rate (43.2 Gbps available throughput capacity), and interconnect B, which has 128 Gbps total throughput capacity and 80% average usage rate (25.6 Gbps available throughput capacity). Here, interconnect A has more available throughput capacity than interconnect B, so interconnect A may be chosen.

Reliability may relate to the level of confidence with which the data storage of the secondary node replicates data in the data storage of the primary node. If a user places emphasis on reliability the secondary node may be chosen to have redundant host bus adaptors and highly reliable, enterprise level storage, such as RAID level **1**. Cost may relate to the cost of using equipment, such as maintenance costs. Cost may also relate to the cost of acquiring currently unavailable equipment. If a user places emphasis on cost, the secondary node may be chosen to have much lower performance than the primary node, in terms of processing capability (server type), storage capability (throughput, cache size, RAID level, etc.), and network interface capability (number and performance of host bus adaptors). For example, storage equipment of RAID level **5** may be chosen.

Thus, by specifying general policies such as (1) performance, (2) reliability, and (3) cost, to follow in establishing

## 11

the clustering system, the user is able control the design of the clustering system, without being required to decipher the detailed considerations relating to technical specifications of related equipment and software. The user may be presented with various general policies from which to choose. The user  
5 may specify policies by simply identifying particular policies as important. The user may also specify policies by assigning importance, or weight, to particular policies. This may be done in different ways, such as by user input of ratings, ratios, percentages, or other measures for different  
10 policies.

The next step under automatic configuration is step **706**, in which information on the current system is gathered. Such information may include the contents of mapping tables, discovery tables, topology tables, and configuration tables.  
15 This information provides a detailed picture of the various aspects of the current system, including the mapping from applications to resources they utilize, available resource and their configurations, and so on.

In step **708**, the information on the current system gathered in step **706** is analyzed to select the most appropriate resources and/or arrangements to be used for creating the secondary node. This is done in view of the various policies, and possibly weights assigned to those policies, as defined  
20 by the user in step **704**. In step **710**, the selected resources and/or arrangements are presented to the user, and the user is given to opportunity to confirm the selection of resources and/or arrangements. If the user confirms the selection, the process continues with step **712**, discussed below. If the user  
25 does not confirm the selection, the process loops back to step **704**.

In step **712**, the selected resources and/or arrangements are used to create the secondary node. If the selected resources need additional software installation or configuration in order to function properly as the secondary node,  
35 such installation or configuration may be performed. Alternatively, the automatic configuration routine or semi-automatic configuration routine may re-select from resources that do not require additional software installation or configuration. Also, default resources that do not require additional  
40 software installation or configuration may also be selected in order to avoid such installation or configuration of software. Finally, in step **714**, the configuration table(s) are updated to include information on the secondary node just created.

## Semi-Automatic Configuration

Under semi-automatic configuration, establishment of a clustering system also begins with step **702**, which has been discussed previously. Next, in step **716**, information on the current system is gathered. This step is similar to step **706**  
50 discussed above. In step **718**, one or more potential selections of appropriate equipment and/or arrangements to be used for creating the secondary node is presented to the user. The user is given the opportunity to select the various equipment and/or arrangements to be used in creating the secondary node. In step **720**, the user's selection is received and presented back to the user for confirmation. Here, a visual topology diagram such as the one shown in FIG. **8** may be presented to the user. FIG. **8** may also represent a  
60 simplified version of block diagram shown in FIG. **1** if the user confirms the selection, the process continues with step **712**, which is has been described previously. If the user does not confirm the selection, the process loops back to step **618**.

In addition, semi-automatic configuration may also take  
65 into account user-defined policies, as is done in the case of automatic configuration. Here, such policies may allow

## 12

potential selections of equipment and/or arrangements presented to be narrowed, so that the user may be presented with a more focused set of potential equipment and/or arrangements from which to make a selection. Other features  
5 discussed above in relation to automatic configuration may be adopted for use with semi-automatic configuration, and vice versa. For example, the visual confirmation diagram discussed in relation to semi-automatic configuration may also be used with automatic configuration, in order to  
10 present the automatically selected equipment and or arrangement to the user for confirmation. Further, variations on the different steps shown in FIG. **7** may also be adopted.

FIG. **1** is a block diagram of a clustering system **100** in accordance with at least one embodiment of the present invention. Such a diagram would allow the user to visually inspect a proposed configuration for a clustering system. This provides an efficient way to present a proposed configuration to the user for confirmation.

Although the present invention has been described in terms of specific embodiments, it should be apparent to those skilled in the art that the scope of the present invention is not limited to the described specific embodiments. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will,  
25 however, be evident that additions, subtractions, substitutions, and other modifications may be made without departing from the broader spirit and scope of the invention as set forth in the claims.

## What is claimed is:

**1.** In a system including a management server, a plurality of processor devices, and a plurality of storage devices, the processor devices and the storage devices being coupled to the management server via a network, a method for the management server to establish a clustering system, said method comprising:

storing storage device information relating to each of the storage devices;

receiving a command from a user interface for specifying at least one candidate of secondary storage devices to which data stored in a primary storage device can be copied and at least one candidate of secondary processor devices, the command having information for identifying a primary processor device and the primary storage device which are configured to be a primary computer resource; and

specifying the at least one candidate of secondary storage devices corresponding to the primary storage device based on the storage device information and the at least one candidate of secondary processor devices corresponding to the primary processor device, the at least one candidate of secondary storage devices and the at least one candidate of secondary processor devices being configured to be at least one secondary computer resource for the primary computer resource,

wherein the storage device information comprises storage capability information including throughput, cache size, and RAID level of the storage devices.

**2.** The method of claim **1** wherein if the primary storage device has storage-based remote mirroring function, the at least one candidate of secondary storage devices is specified from the plurality of storage devices, which has storage-based remote mirroring function and is accessible from the primary storage device.

**13**

3. The method of claim 1 wherein the primary computer resource includes a first network interface device for the network; and

wherein said at least one secondary computer resource includes a second network interface device for the network. 5

4. The method of claim 1, wherein said at least one candidate of secondary processor devices is selected from

**14**

the plurality of processor devices having processing capability equal to the primary processor.

5. The method of claim 1, wherein the user interface includes GUI.

\* \* \* \* \*